

REMARKS

Claims 8-10 are pending.

Claim Rejections Under 35 U.S.C. §102

Claim 8 is rejected under 35 U.S.C. §102(b) as being anticipated by U.S. Patent No. 5,418,916 to Hall et al. ("Hall").

Hall discloses a multi-processor system that detects errors during processors execution of instructions and returns to a previous point in the execution stream to re-execute the segment that led to the errors. Along the execution stream (consisting of executing instructions down the stream), checkpoints are set up at selected instruction stream intervals. At each checkpoint, all relevant data for each instruction are retained so that the execution can be restarted from each checkpoint. This is facilitated by a store buffer or cache that stores the data during the checkpoint period. The relevant data for a given checkpoint is retained in the store buffer until the end of the given checkpoint period and is then transferred into slower memory such as L2 or main memory.

"For all central processing unit (CPU) initiated storage operations following a checkpoint, the modified data cache pages are saved in the stored buffers until the storage operations are completed for the checkpointed instructions. In the case of the processors with a store-in level two (L2) cache design, the modified cache data is saved in the level one (L1) store buffer until the checkpointed instructions have been completed. It is then written to the L2 cache and become visible to other processors. For processors with a store-through cache design, the modified cache data is also buffered in the L1 store buffer. When a checkpoint period is completed, the saved cache data is stored in memory." (Col. 4, lines 2 to 15.)

Since the data is backed up in the context of a processor executing instructions, errors detected are corrected as soon as possible. This accomplished by rolling back to the beginning of the checkpoint period and using the retained data in the store buffer to restart the instructions since the beginning of the checkpoint period. At any time, only one checkpoint period need be maintained. At the end of the period, the checkpoint and its data in the store buffer are discarded and are replaced by an updated checkpoint. For example see col. 10, lines 10-21:

“The global controls keep track of all GPR operations and detect the completion of all EPU operations within a checkpoint. Thus, the EPU need not actively signal completion of its operations. When the IPPU detects that all of the processing complete bits are set, the IPPU begins to clear that checkpoint. A checkpoint is cleared when the backups are invalidated and the PSW stack entry is marked complete. The IPPU invalidates the backups associated with the checkpoint by broadcasting to the processing units. The PSW entry is marked complete at the same time. After the backup invalidation and PSW entry completion, all information associated with the checkpoint becomes invalid and the backup arrays and PSW entry can be used for a new checkpoint.”

Thus, unlike the present claimed invention, the issue of storing a time series of checkpoints and retaining all their associated data never arises.

The current invention is intended for backup of data that a user may want to roll back to any number of a series of snapshots. Thus, a time series of snapshots is maintained at any one time and each of their associated data must be recoverable. As pointed out in the Specification, conventional snapshot systems maintain each snapshot relative to the current state of the system. While the restoration is more convenient, involving only the snapshot in question and the current state, the associated data necessarily contain a fair amount of redundancy. (See for example, FIGURE 4 as compared to FIGURE 3 of the Specification.) The claimed invention retains a minimum of data necessary for maintaining a time series of (“multiple”) snapshots with substantially little or no redundancy of preserved data between the series of snapshots.

Independent claim 8 is being amended to more clearly claim the invention. The “multiple” nature of the snapshots and the lack of “redundancy of preserved data” are being explicitly recited as claimed elements. The feature “terminating the cache associated with the current snapshot just prior to the start of a next snapshot” is disclosed in box 540 of FIGURE 5. (See also the first part of [0077].) Applicants maintain that these features are not taught or suggested by Hall’s one-trackpoint-at-a-time system.

Accordingly, claim 8 is believed allowable over Hall. Reconsideration of the rejection is respectfully requested.

Claim Rejections Under 35 U.S.C. §103

Claims 9 and 10 are rejected under 35 U.S.C. §103(a) as being unpatentable over Hall.

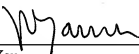
Claims 9 and 10 are dependent on a believed allowable claim 8, and are therefore allowable too. Reconsideration of the rejection is respectfully requested.

Conclusion

Accordingly, it is believed that this application is now in condition for allowance and an early indication of its allowance is solicited. However, if the Examiner has any further matters that need to be resolved, a telephone call to the undersigned at 415-318-1165 would be appreciated.

FILED VIA EFS

Respectfully submitted,


Philip Yau
Reg. No. 32,892

October 4, 2006
Date

PARSONS HSUE & DE RUNTZ LLP
595 Market Street, Suite 1900
San Francisco, CA 94105
(415) 318-1160 (main)
(415) 318-1165 (direct)
(415) 693-0194 (fax)